arXiv:1603.08059v1 [cs.PL] 26 Mar 2016

CVC Verilog Compiler – Fast Complex Language Compilers Can be Simple

Steven Meyer

Tachyon Design Automation
Boston, MA 02109, USA
smeyer@tachyon-da.com

Abstract

This paper explains why the CVC Verilog hardware description language (HDL) optimized flow graph compiled simulator is fast. CVC is arguably the fastest full IEEE 1364 2005 standard compiled Verilog simulator available yet consists of only 95,000 lines of C code and was developed by only two people. The paper explains how CVC validates the anti-formalism computer science methodology best expressed by Peter Naur's datalogy and provides specific guidelines for applying the method. CVC development history from a slow interpreter into a fast flow graph based machine code compiled simulator is described. The failure of initial efforts that tried to convert CVC into interpreted execution of possibly auto generated virtual machines is discussed. The paper presents evidence showing CVC's speed by comparing CVC against the open source Icarus simulator. CVC is normally 35 to 45 times faster than Icarus, but can be as fast as 111 times or as slow as 30 times. The paper then criticizes the competing Allen and Kennedy theory from their "Optimizing Compilers" book that argues fast Verilog simulation requires detail removing high level abstraction. CVC speed comes from efficient low level usage of microprocessor instruction parallelism. The paper concludes with a discussion of why the author believes special purpose full accurate delay 1364 standard hardware Verilog simulators and parallel Verilog simulation distributed over many processors can not be faster. CVC is available as open source software.

1. Introduction

This paper presents a novel method for implementing fast compilers for complex computer languages. The simple organization and development methods used to create a fast Verilog hardware description language (HDL) machine code simulator are described. The organization of CVC can be viewed as a modern analog of the multi-pass compilers applying problem specific methods used to develop early compilers (Gries 1971) (8-10 for history, 410, 411 and 451-454 for methods). The early compilers were developed by scientists who were mostly trained as physicists. See Budiansky's account of the role of English physicist Patrick Blackett's department during WWII that was perhaps the first use of modern algorithmic thinking (Budiansky 2013).

Computers are now so fast and are equipped with so much fast random access memory that the various (virtual) passes can be executed without storing results of each pass in secondary storage. The organization in which one unified representation is continually modified and transformed is similar to the approach Peter Naur used in developing the Gier Algol compiler (Gries 1971) (p. 9). For example, pass 6 for a compiler run on a machine with only 128,000 words of memory checked types of identifiers and operands (and updated information) then converted to Polish notation and output the result for the next code generation pass. In CVC the code generator does similar checking but converts the source program (really the internal representation of source because of ample memory) to a flow graph of basic blocks that contain virtual machine instructions.

2. What is CVC

CVC is an electronic design automation (EDA) simulator for models of electronic hardware described in the IEEE 1364 2005 Verilog HDL standard (IEEE-Standards-Board 2005). CVC implements the entire standard including both register transfer level (RTL) procedural simulation and accurate delay annotated delay gate level simulation. CVC is marketed using the open source business model. It is available at *http://www.tachyon-da.com* by clicking on the **Sign up and Download** button (Tachyon-Design-Automation-Corp. 2015)

CVC was developed by two people in less then 10 years. The ten years included training a young programmer who started as a college student intern and includes most of the author's time implementing the elaboration code for the generate feature added by the 1364 standard committee which is almost incompatible with CVC's instance based display machine register technology (**__idp** area).

CVC consists of about 330,000 lines of C code that is 10% or less the size of the competing commercial quality Verilog full accuracy simulators. CVC was first developed as an interpreted Verilog simulator then the flow graph based compiler was added. The Verilog elaborator (parser, fix up and simulation preparation phases) is about 100,000 lines (numbers are approximate because there are common routines and overlap). Of the 100,000 lines, 20,000 or 20% are needed for the complicated Verilog 2005 generate feature.

Generate allows compile time variables called parameters to change HDL variable sizes and design instance hierarchy structure. This feature creates instance specific constants that must be treated as variables during simulation. The other Verilog simulators flatten designs resulting in much larger memory use for designs with many repeated instances. It is not uncommon to have millions

[Copyright notice will appear here once 'preprint' option is removed.]

of instances of a latch or flip flop macro cell each of which requires storing not just its per instance state information but also the machine instructions for every instance. CVC stores one instance model plus per instance state information. The development of the Verilog 1364 standard by committee has continually made Verilog more complicated (the language reference manual is 590 pages) and more difficult to implement CVC's instance base pointer algorithms. CVC's display offset algorithm results in a simpler code generator and faster simulation at the cost of a much more complicated elaboration phase.

About 50,000 lines are needed to execute interpreted simulation. The compiler itself is only about 95,000 of which 15,000 lines are executable binary support libraries. 70,000 lines are needed to implement miscellaneous features used both by the compiler and the interpreter: SDF delay annotation, four programming language interface APIs (tf_, acc_, vpi_ and dpi_), a debugger for the interpreter, toggle coverage recording and report generation, rarely used switch level simulation, plus an expression evaluation variant called X propagation that implemented a more pessimistic unknown (X) injection algorithm.

3. Relation to Naur's Datalogy and Anti-Formalism

The theoretical background behind CVC's development method follows Naur's methods. In the 1990s, Peter Naur, one of the founders of computer science realized that CS had become formal mathematics separated from reality. Naur advocates the importance of programmer specific program development that does not use preconceptions. The clearest explanation for Naur's method that was used in developing CVC appears in the book *Conversations - Pluralism in Software Engineering* (Naur 2011). This books amplifies the program development method Naur described in his 2005 Turing Award lecture (Naur 2007). In (Naur 2011) page 30, the interviewer asks "... you basically say that there are no foundations, there is no such thing as computer science, and we must not formalize for the sake of formalization alone". Naur answers, "I am not sure I see it this way. I see these techniques as tools which are applicable in some cases, but which definitely are not basic in any sense." Naur continues (p. 44) "The programmer has to realize what these alternatives are and then choose the one that suits his understanding best. This has nothing to do with formal proofs." Einstein described this as the 20th century split between axiomatics and reality by saying that "axiomatics purges mathematics of all extraneous elements" which makes it evident "that mathematics as such can not predict anything" about reality (Einstein 1921). See also (Naur 2005) and my 2013 IACAP paper (Meyer 2013) for more detailed discussion of Naur's anti-formalism.

Current compiler development methodology chooses formal algorithms over Naur's programmer specific approach that rejects pre-suppositions. During the development of CVC, I was also aware of anomalies in mathematical foundations of logic that current computer science takes as truth beyond criticism. The first is Paul Finsler's proof that the continuum hypothesis is true (Finsler 1969). The proof is only indirectly related to computer science. The second example is Juri Hartmanis' proof that P=NP in PRAM (parallel ram) models (Hartmanis and Simon 1976). A problem shift then occurred which used PRAM models only for studying concrete algorithm complexity but kept Turing machine models for studying abstract complexity. This power of PRAM machines contributed to looking for other sources of parallel speed improvement in CVC.

A more direct result of rejecting formalism is the choice to use the fast Cooper-Kennedy dominator algorithm (Cooper et al. 2006). This choice and its use in the crucial optimization algorithm define-use lists is easy once one starts with skepticism toward algorithms that have been formally proven to be correct. I see the Cooper algorithm's "real" speed as falsification of the axioms used in concrete complexity theory.

The main lesson of CVC's anti-formalism is to avoid using abstraction. Avoid machine generated tables and compiler phases generated by automatic generators. Parse using language specific recursive descend and use the C run time stack for remembering context. Parse expressions using simple operator precedence. This results in a very fast parser in CVC but is only possible because assignment operators are not part of expressions in Verilog. Assume proofs use axioms that do not apply to reality unless specifically determining that the axioms are good. Finally, following Naur, look for algorithms that are simple so they can be adopted to the problem specific aspects of Verilog and the previous organization and algorithms of CVC.

4. CVC Development History

CVC (called Cver at the time) was originally developed in the 1990s as a Verilog simulator to compete with the original Verilog XL simulator that used interpreted execution. At the time Verilog semantics was aimed at interpreted simulation because simulation properties such as delays could be set at any time during a simulation run and because a command line debugger was an integral part of Verilog, i.e. simulations often expected to read Verilog source from script files at various times during a simulation run. In the late 1990s Verilog native machine code compilers were introduced and Verilog semantics was changed to require specification of simulation properties at compile time. See (Thomas and Moorby 2002) for a historical description of the Verilog HDL.

By 2000 CVC was no longer speed competitive so it was used as the digital engine for the Antrim Verilog AMS (analog and mixed signal) simulator. The CVC speed problem was not as serious for mixed signal simulation because analog simulation that requires solving differential equations runs orders of magnitude slower than digital simulation. The Antrim AMS project allowed CVC to be improved from use by the most sophisticated electronics companies.

After the end of the Antrim project, CVC needed to be improved as a digital simulator that would be speed competitive with compiled Verilog simulators. The most obvious speed problem involved register transfer level (RTL) simulation. RTL simulation is almost the same as normal programming language execution except Verilog values require at least 2 computer bits per Verilog bit and the RTL execution must interact with an event driven scheduler.

The most obvious problem was that a number of **if** statements in the interpreter C code were needed to select which interpreter algorithm to run. For example, a simple logic and (&) operator needed different evaluation C language code sections (usually procs) for scalars (1 bit), narrow vectors (up to 32 or 64 bits), wide vectors and strength model bit vectors (one byte per bit required). In addition for all but scalars, both unsigned and signed cases were needed. Sign extension for non integral number of word bit vectors requires significant calculation. The Verilog standard requires that at least up to one million bits wide vectors simulate correctly. It seemed that taking the interpreter evaluation routines and converting to high level instructions for a virtual machine that could then be interpreted was a good idea (see for example (Ertl and Gregg 2003)). Also, automatically generating interpreter code was tried (Ertl et al. 2002). The development was not too hard, but the resulting execution speed increased performance by only a small amount. Although the **if** statement overhead was removed, extra overhead to decode and execute the interpreter virtual instructions nullified most of the gains.

4.1 Concrete algorithms but not organization from Morgan's optimizing compiler book

It was realized that the instruction level parallelism and branch prediction in modern microprocessors units was needed. Development of a full code generator began. We next attempted to implement the concrete step by step method in Robert Morgan's book on building an optimizing compiler (Morgan 1998). The book describes the method used in building the very good Digital Equipment Corporation Alpha microprocessor compilers.

CVC does not use the code generator organization from (Morgan 1998) (section 2.1, 21-26). Morgan advocates a basically breadth first filter down approach with a chain of transformations each of which has a different data representation. Morgan writes: "Each phase has a simple interface" that can be tested in isolation. Also, "No component of the compiler can use information about how another component of the compiler is implemented" (p. 21).

Instead, CVC implements modified versions of the very good algorithms spread throughout the Morgan book. Also, during development of the CVC code generator, our new ideas were compared against the concrete Morgan book approach to make sure they were no worse. The Morgan book algorithms are especially useful because exceptions are discussed. For example allowing non static single assignment (SSA) form exceptions that violate the rule that each variable is assigned to only once or sometimes that define-use chains do not have exactly one element (p. 142).

CVC uses Morgan's idea that virtual instructions should be as close to machine instructions as possible (Figure 2.2, p. 24). The exception in CVC is that Verilog requires very complicated mostly boiler plate prologue and epilogue instruction sequences. Those higher level virtual instructions are modified versions of virtual instructions from the original interpreter.

CVC code generator flow uses one master representation accessible from the interpreted data base. Both flow graphs and temp names (unbounded number of virtual registers that in Verilog are often wide and contain two components to represent 4 values) are accessed in numerous ways. Basic blocks are accessible directly from interpreter execution form that then points to flow graphs. Flow graphs especially for net change propagation operators are accessed from indexed tables and AVL trees (see **igen.h** in CVC source).

I believe the CVC organization that combines all information into one master data base is better. Any code generation phase can use any of the information that is accessed either directly from the interpreter execution net list, from indices (sometimes indexed tables and sometimes trees), or from code generation records. Idea is to continue to improve the global data base during code generator development and to continually add more information to all of the various parts of the one unified representation. For example, flow graph building algorithms were used to improve design elaboration data structures and interpreter execution data structures.

This unified data base where each part is kept consistent with other parts is the key to CVC's simplicity and code quality. For Verilog, because there are so many different types of operations from procedural RTL to declarative gate level to load and driver propagation, depth first code generation is better. Morgan's type of low level machine instructions (p. 24) are generated with some optimization by expanding constructs all the way down to something close to the final virtual instruction sequences when the flow graphs virtual instruction sequences are built. Later mapping to machine instructions except for X86 fixed registers is straight forward. This approach may be Verilog specific because Verilog allows values to be read and written from anywhere in a design because of cross module references and the programming language interface (PLI) can run in any Verilog thread. There is effectively no usable context information in Verilog.

For example, once some C code for the unified data structures of the flow graph and basic block mechanism were written, the Morgan book detailed algorithms on define-use lists and importance of SSA (12.5.1 p. 291 and 7.1, p. 142) could be applied. CVC uses the heuristic to generate (top down depth first) lots of temporaries and then fixing SSA problems when needed or even allow some constructs to violate SSA form. The crucial data structure used for optimization is the define-use lists. Morgan suggests the idea of allowing non SSA instructions and temporaries (p. 142). Extra temporaries can then be eliminated during optimization passes through flow graphs and virtual instruction lists. CVC uses an **I_COPY** virtual instruction because of the large number of different data representations in Verilog.

4.2 CVC file organization

Basic block creation and virtual instruction generation C code is in CVC's **v_bbgen**'s C files. Define-use list and other flow graph elaboration code are in the **v_bbopt.c** file. The **v_regasn.c** file assigns machine registers to the unbounded number of temps. The **v_cvcms.c** and **v_cvcrt.c** files plus the **v_asmlnk.c** file contain support C procs plus code to generate the GNU AS assembly output, run **gas** and link the final output executable **cvcsim**. The **v_aslib.c** file contains wrapper C procs that are called from the generated assembly but whose function is to call an interpreter execution proc. By using wrappers, early versions of the compiler could compile almost all of Verilog but simulation was not yet fast because execution used wrappers that just executed the slow interpreter code.

5. What is Verilog

Verilog is a Pascal like language (Wirth 1975) for the description of electronic hardware. All variables are static because there are no implicit stacks in hardware. Verilog is a combination of normal behavioral programming with parallelism, execution of hardware described at the RTL level and low level primitive declarative gates and flip flops. A common design method is to code circuit descriptions in RTL then run a program to synthesize the RTL into gates also coded in the Verilog language. Verilog is used to simulate (predict behavior when an IC is fabricated) both the RTL and the synthesized gates with accurate timing. See (Thomas and Moorby 2002) for a description of the Verilog HDL. See (Sutherland et al. 2006), pp. 401-413 for a history of the Verilog HDL. See (Allen and Kennedy 2002), pp. 619-622 for a description of Verilog from the viewpoint of optimizing compiler development.

6. Why Verilog Simulations Run Slowly Compared to Computer Programs

First, Verilog RTL values require 2 bits (4 values) for every hardware bit. Gate level accurate delay simulation also requires bus values which have 127 different values and driving strengths. A simple logic operation requires at least 3 or 4 instructions (not counting loads and stores). Second, hardware registers are almost always wider than the native machine register width because hardware design involves modeling the next generation electronics. This requires evaluating multiple machine words for each operation. Third, Verilog requires event driven simulation. When a delay or event control (@(clk3) say) is executed, the simulation must schedule a new event in an event queue and suspend the current execution thread to be restarted later. A even slower process occurs when a value is changed. All variables that are on right hand side expressions driving the value must be evaluated and new assignments made. These effected variables are called loads. Also, when a wire is evaluated it may be necessary to evaluate multiple drivers

and determine which is the strongest. See (Meyer 1988) for a data structure that allows implementation of efficient load propagation and driver competition algorithms.

7. CVC Performance

CVC is arguably the fastest full accuracy 1364 Verilog simulator. Evidence here is anecdotal from users because the other commercial quality simulators are licensed with no public benchmarking allowed clauses. There is one anonymous published comparison of CVC RTL but not gate level simulation speed (Anonymous 2011). Measurements showed CVC to be 2.8 times and 5 times faster than a commercial Verilog simulator for the two discussed designs.

A source of anecdotal evidence is responses from people who run some open core designs that are included in the CVC release "tests_and_examples" "/open_cores" directory. CVC optimization is so good that in one case full accuracy 1364 Verilog simulation is almost as fast as a cycle based simulation that omits nearly all details because the model can be simulated with the +2state option (see the m68k model in the CVC releases' open_core directory). CVC quality has actually been validated by a ruling in a US District Court case (US-Federal-Court 2003).

7.1 Comparison with Open Source Icarus Verilog

There is another open source Verilog simulator Icarus Verilog called **iverilog** that implements most IEEE 1364 Verilog RTL features. It is difficult to use for benchmarking because it does not implement Verilog source macro cell library options (called **-v and -y** libraries). The CVC developers simplified some publicly available small open core RTL designs (ranging from 500 to 14,000 lines of Verilog source) so that macro cell source libraries are not required. The examples with canned Verilog file lists are in the open source CVC release directory **open_cores** that is in the **tests_and_examples** directory.

Sixty four bit CVC64 is usually 35 to 45 times faster than Icarus Verilog release called **10.0 stable** made with the default options (-g -O) run on an Intel Core i5-5200U 2.2 GHz low power processor running Linux Centos 6.7. There are two outliers. An open core model of an old Motorola 68K processor is only 32 times faster. CVC simulation runs in 15.9 seconds versus Icarus 522.0 seconds. The model uses continuous assigns requiring event scheduling. The event queue can not be compiled. CVC has an optimize zero delay non blocking assignment option (+nabopt). When it is used to simulate the m68k core, CVC is 35 times faster. At the other extreme, a SHA1 hash function calculation RTL model runs 73 times faster with CVC. CVC simulation takes 1.57 seconds and with the +2state option 1.03 seconds. Icarus takes 115 seconds. If the CVC two state simulation option (+2state) is used, CVC is 111 times faster. sha1 is a C program translated to Verilog that has almost no delay scheduling. A more average open core such as ac97 that is an audio coder simulates in 79 seconds (70 seconds if the +nbaopt and +2state options are both used) on CVC and 3,016 seconds on Icarus. CVC is 38 and 43 times faster.

The CVC two state option works by keeping the X and Z, B part words around but only needs to initialize the the words to zero once. If a design simulation really requires X and Z values, simulations run with the option will be incorrect. CVC flow graph optimizer normally removes B part basic blocks when B part X and Z values can not occur even if the two state option is not used. The option allows generating simpler flow graphs that allow evaluation of value A parts to be optimized more.

8. Simple Design Method for Complex Language Compilers using CVC Example

8.1 Develop an interpreter using language specific simple concrete methods

CVC was developed at the same time the Verilog IEEE 1364 Verilog standard was developed and was continually changing. Also, CVC was developed in a period when simulation needed to match the original Gateway Design (then Cadence Design) results. There is no way to simplify or shorten this process. Some helpful ideas are:

8.2 Use simple organization

For CVC one centralized include file is used by every C source file with defined and used prototypes at the top of each source file. This organization makes it easy to eliminate occurrences of more than one routine for the same basic function (wide vector sign extension for example). Also, it allows making a change in only one place when Verilog changes.

8.3 Use only one internal organization (net list data structure for Verilog)

CVC scans and parses Verilog source into normal module lists, statement and gate lists, expression trees and linked symbol tables. Then the next set of fix up procedures is used to fill the same date structure with more information which includes possibly totally changing the instance tree hierarchy. Then the next phase fills in more simulation preparation information and allocates variable and state memory just before beginning simulation. This organization allows easily moving processing steps forward and backwards when the language changes.

8.4 Avoid generators, grammars and tables

The simplest and most powerful scanning method is to use a giant case statement. The simplest and most powerful parsing method is to use language specific recursive descent. This is especially important in complicated languages such as Verilog where the scanner needs parser information and the parser needs scanner information. The 1364 standard committee does not worry about context free grammars. Verilog assignments are separate from expressions so CVC is able to use simple and fast operator precedence parsing (Gries 1971) (section 6.1, 122-132). CVC can elaborate 5 million line designs is less than 15 or 20 seconds on a modern fast CPU. Fast elaboration speeds up development of code generator algorithms, allows faster compiler debugging and assists in finding faster machine code patterns to generate.

8.5 For complex languages first develop an interpreter

The first step in fast compiled Verilog development is to get an exact interpreted simulator so that compiler machine code can be regressed against the interpreter standard. Once simulation is running, many speed improvements will become obvious (made visible) without needing to wait for a long design phase before experimental data is available.

8.6 Add interpreter wrapper capability

CVC defines **I_CALL_ASLPROC** and **I_CALL_ASLFUNC** virtual instructions that works with the normal unbounded register temps and define-use predominator algorithms, i.e. interface is no different from a low level **I_MOV** virtual machine instruction. This allows running simulations with only some constructs compiled. In the early development of CVC, only narrow (less than 32 or 64 bit) variables were compiled. All wider expressions and assignments were evaluated with wrappers. Then step by step wrappers can easily be replaced by generated low level machine instructions. Some

very complicated Verilog algorithms such as multiple path delay selection and switch level simulation are still just wrappers.

8.7 Make writing virtual instruction generation same as interpreter execution

In CVC something like binary operator evaluation starts with a wrapper, then is replaced by a version of the same proc with calls to **__gen_tn** to create temporary registers, and **__start_bblk** to replace **if** statements (see **eval_binary** proc in **v_ex2.c** file near line 6600) that is in file **v_bbgen.c**. This allows code generation to be mostly mindless recoding of interpreter evaluation into basic block temp, basic block and virtual instruction generate proc calls. Code generation for more complicated simulation operations such as event processing can be simplified this way also.

8.8 Modify Morgan's dominator-based optimization algorithms

The book *Building an Optimizing Compiler* by Robert Morgan contains a concrete simplified method for optimizing flow graphs and computing define-use data structures. Start with that and simplify even more and modify to fit your complex language. This is the hard very important part of developing a fast compiler. Once good flow graph optimizations are implemented, register allocator becomes much easier and better. See **__optimize_1mod_flowgraphs** at the beginning of file **v_bbopt.c** for the code that implements this.

8.9 Use low level virtual instructions close to machine instructions

CVC defines and emits virtual instructions that are mapped into machine instructions. It is best to generate low level virtual instructions because the code generation human coder for the particular Verilog feature can craft good instructions sequences. Verilog operations at minimum require operating on an A part for the 0 and 1 value and a B part for the X and Z value so even simple operations require multiple machine instructions and probably two by two vector cross product evaluation.

The copy operation is an exception especially for Verilog because much simulation is copying data that can be very wide. Original code generation inserts a complicated virtual **I_COPY** instruction whenever there is a possibility of a need for a copy. Then during mapping from the low level **I_COPY** to machine instructions, the copy usually can be removed without needing to emit anything.

8.10 Avoid extra representations such as tuples and breadth first transformations

One of the best simplifying and simulation speed increasing methods in CVC is depth first code generation. This method is intentionally the opposite from what (Morgan 1998), page 212 recommends. Morgan recommends breath first code generation with successive transformations to lower levels. It is much simpler to use depth first instruction generation of almost machine instructions because it allow the compiler writer to control machine code sequences and reduces number of lines of compiler code.

8.11 Use experimentation to find fast CPU instruction parallelism code sequences

Once the CVC code generator was written and debugged, the best speed improvement method was to set up shell scripts that would run two different versions of the compiler on a speed test regression suite and compare results. The low level machine instruction sequence for the best would then be used. We did not have access to the X86_64 multi-issue pipeline optimization rules documentation so we needed to run experiments. Maybe the experimental approach is better in general because the rules are so complex. This was the most important simulation execution speed up idea. The other was compiling the change propagation and scheduling event processing code into flow graphs.

9. Allen and Kennedy Hardware Simulation as Abstraction Contradicted by CVC

In the book *Optimizing Compilers for Modern Architectures*, Randy Allen and Ken Kennedy argue that the task of optimization of hardware descriptions is to abstract to a less detailed level (Allen and Kennedy 2002). Allen and Kennedy write "Another way of saying this is that simulation speed is related to the level of abstraction of the simulated design more than anything else – the higher the level, the faster the simulation" (p. 624). CVC has shown that low level exact detail modeling is faster because it allows maximum low level instruction parallelism that is built into modern microprocessors. This parallelism can be utilized by the compiler's code generator for faster execution. The optimizations include processing as wide a bit vector chunk at a time as possible, generate instruction sequences to keep instruction pre fetch and pipe lines full and generate instruction sequences that work well with branch prediction algorithms. Basic blocks in Verilog are small but very numerous because separate A parts and B parts require conditionals.

CVC shows that not very complicated code generator (less than 100,000 lines of C code) can produce good code using the experimental method of run speed regression test suites with different version of generated instruction sequences and choose the fastest. The standard way (almost but not quite required by the IEEE P1364 standard) for storing Verilog 4 value logic and wire values is to store values as two separate machine words. The B part selects unknown values: X (unknown) and Z (high impedance) values. If the B part is zero, the A part selects 0 or 1 values. Bit and part select operations need to be optimized as fast load, shift and mask operations that use mostly boiler plate instruction sequences. Four value logic operations can also be generated as combinations of machine logic instructions combining the A and B parts. Although, complex logic operations are sometimes better with A and B parts separated or even as table look ups.

The remainder of this section criticizes specific ideas for implementing simulation abstraction from the Allen and Kennedy book.

9.1 Inlining modules, p. 624

Expanding modules inline is called a fundamental optimization. CVC disagrees because it is much better to use the normal computer language optimization of not inlining but separating procedure code form state and variable information with based addressing using traditional display technology (Gries 1971) (pp. 172-175). Not only is code smaller but separating state from model description makes many optimizations visible especially ideas for for pre-compiling event processing.

9.2 HDL level execution ordering, p. 626

Allen and Kennedy argue for HDL statement reordering optimizations. Far more important is machine instruction reordering to maximize low level microprocessor instruction parallelism. Part of the reason for this is that although Verilog has fork-join constructs, they are rarely used. Parallelism in Verilog comes from a very large number of different always block usually triggered with an event control (for example **always @(clk)**).

9.3 Dynamic versus static scheduling, p. 627

This section discusses oblivious evaluation. Instead of the normal Verilog algorithm that "dynamically tracks changes and propagates those changes", the alternative is too blindly evaluate without any change propagation overhead. The oblivious method can not work

for Verilog because it is common to have thousands of tick gaps between edges that have very high event rates. In CVC static scheduling and compiling event after change recording into pre-compiled flow graphs as much as possible is more important.

9.4 Fusing Always Blocks, p. 628

Here Allen and Kennedy are correct. Fusing always blocks makes a huge speed improvement.

9.5 Vectorizing Always Blocks, p. 632

The idea is that Verilog code generation optimization should "rederive the higher-level abstraction that was the original intent". The problem with this is that the decision to code scalarized or vectored is best left to the designer (HDL generation program). Vectorizing is not always better because if only a few bits change per clock cycle in a wide vector, it is better to simulation the scalarized individual bits. There is no way to determine switching frequency for a given simulation without running it. Change detection of vector selects is fast because it only requires a mask, shift and xor then branch, but propagation of the changes can be expensive if the selected regions used in right hand side expressions do not exactly match. It is much better not to change the Verilog HDL code, but to find fast instruction sequences.

9.6 Two-State versus Four-State Logic, p. 637

Two state simulation is good when it is possible. The problem is that the point to hardware simulation is to find mistakes that will result in unknown X (IC state will be sometimes 1 and sometimes 0) states in the fabricated integrated circuit. Most real HDL designs do not allow two-state simulation, but if possible simulation is much faster because evaluations can directly use hardware machine instructions. Most four state evaluations are similar to evaluating two by two vector cross products. The low level design feature in CVC is to treat temporaries even for four state values that require an A part word section and a B part word section as one temp. When an expression evaluation can be executed as two state, the B part instruction sequences are just not emitted. For Verilog, variable can be declared to be two-state, but in mixed two-state and four-state evaluation, the two-state values must be treated as four-state with zero B part.

9.7 Rewriting Block Conditions, p. 637

Allen and Kennedy advocate changing trigger conditions on Verilog blocks by abstracting and guessing user intent to eliminate the need to propagate change operators that in Verilog are usually event controls on always blocks. Instead of trying to rewrite the Verilog source to higher level, it is better to in line event queue propagation by generating flow graphs because most event controls are simple. Change operator occurrence scheduling and wake up event processing are inlined into flow graphs in CVC. In CVC, the state data is first coded into the per instance (**_idp**) based display area (see proc **alloc_fill_ctevtab_idp_map_els** near line 3820 of file **v_prp.c**). Then separate flow graphs are coded to schedule the changes and to propagate the changes. The generated flow graphs are optimized the normal way so most change operators require just a few instructions to schedule and a few instructions when the flow graph is jumped to from the event queue processing code. The scheduling part for delay controls (**@(clk)** say) flow graph generation proc is in the **__gen_dce_schd_tev** routine near line 10400 in file **v_bbgen3.c**.

10. GPU or Other Multi-Core Parallelism Does Not Work for Verilog

There have been many projects that attempt to use computer hardware rather than optimizing flow graph compilers to speed up Verilog simulation. The reason for this is that Verilog simulations are a significant consumer of electronics company compute cycles with sometimes entire server farms dedicated to running only Verilog. Verilog compiled executable speed is of crucial importance and has huge economic value. At least so far efforts to speed up HDL simulation either with special purpose hardware or GPUs have failed. It is possible to build special purpose hardware that will run unoptimized Verilog 10 times faster. The problem is that a good flow graph based optimizing compiler run on a modern fast multi-issue microprocessor can be 10 times or more faster than the naive unoptimized code the special hardware is compared to.

Hardware emulation uses a different approach in which a design converted to gate level is "layed out". Emulated designs that can run at a few hundred thousand cycles per second are good for early software development, but does not "simulate" a design in sufficient detail for debugging.

There have also been projects to simulation Verilog faster using large scale processor level parallelism sometimes using special purpose hardware and sometimes using multiple cores or GPUs. At least so far, all of these projects have failed for full accuracy 1364 Verilog simulation. The reason is that Verilog HDL basic blocks tend to by very small with a high proportion of jumps and synchronizations. Small basic blocks parallelize efficiently on multi-issue CPUs but has too much synchronization overhead for coarser parallelism. The one place CVC uses parallelism (up to 2 additional cores) is to encode value change dump files (FST format is smallest and fastest). This use of parallelism can improve simulation times for simulations that generate huge value change dump files by a factor of two. The very large change value files are written because they can be input into back end physical design tools that use the value changes to optimize IC power usage.

11. Conclusions

The CVC project has been successful in the sense that a fast commercial quality compiler was developed without many resources. CVC is not widely used because the larger Electronic Design Automation industry companies are vertically integrated so customers are to a large extent locked in to one companies' design flow (connection to back end physical design tools). Most CVC customers have unique ways of simulating their electronic designs that usually involve extensive use of the Verilog PLI.

The general area of how to represent digital electronic circuit is currently rather unsettled. Many designers would prefer to write programming language (C usually) code and have a program automatically convert the code into a hardware design that needs to be represented in Verilog (called high level synthesis). One current alternative is the SystemC language (IEEE-Standards-Board 2014), but I think because it does not code event controls in a hardware like manner, it is not used much.

If a design is really just a computer program and implemented as an FPGA, it may be cheaper and better in speed and lower in power to not design hardware but instead implement the device as highly optimized programs on off the shelf multicore SoCs containing conventional CPUs, signal processing CPUs, GPUs and other types of CPUs. Another way of expressing this observation is that if digital circuits are designed using only two state logic with X and Z cross products ignored, implementation as computer programs may be better.

In my view there is a need for a simple Verilog like HDL that is intended to be generated by computer programs - maybe V–. Verilog elaboration and use of constant parameters that are sometimes not really constant at run time makes sense when HDL designs are coded by hand because it makes Verilog easy to write, but not when Verilog is machine generated. If such a language eliminated non locality such as cross module references and force-release state-

ments (originally intended to allow coding reset buttons), much faster simulation might be possible. Such simulation could maybe use some type of graph theory connectivity.

Acknowledgments

Andrew Vanvick(AIV in source) was the other CVC developer.

References

A. Allen and K. Kennedy. *Optimizing Compilers for Modern Archictectures*. Academic Press, 2002.

Anonymous. A verilog simulator comparision, September 2011. URL `http://www.semiwiki.com/forum/content/777-verilog-simulator-comparison.html`.

S. Budiansky. *Blackett's War: The Men Who Defeated the Nazi U-Boats and Brought Sciene to the Art of Warfare*. Knopf, 2013.

K. Cooper, H. T., and K. Kennedy. A simple, fast dominance algorithm. Technical Report TR-06-33870, Rice University Department of Computer Science, 2006.

A. Einstein. Geometry and experience, January 1921. URL `http://www.relativitycalculator.com/pdfs/einstein_geometry_and_experience_1921.pdf`. Lecture before Prussian Academy of Sciences, Berlin.

M. Ertl and D. Gregg. The structure and performance of efficient interpreters. *Journal of Instruction-Level Parallelism*, 5, 2003. electronic archive journal.

M. Ertl, D. Gregg, A. Krall, and N. Paysan. vmgen – a generator of efficient virtual machine interpreters. *Software Practive and Experience*, 32(3): 265–295, 2002.

P. Finsler. Uber die unabhängigkeit der continuums hypothese. *Dialectica*, 23(1):67–78, 1969.

D. Gries. *Compiler Construction for Digital Computers*. John Wiley and Sons, 1971.

J. Hartmanis and J. Simon. On the structure of feasible computations. In M. Rubinoff and M. Yovits, editors, *Advances in Computers*, volume 14, pages 1–43. Academic Press, 1976.

IEEE-Standards-Board. Ieee std 1364-2005 verilog hardware description language reference manual, 2005.

IEEE-Standards-Board. Ieee std 1666-2014 core system c language version 2.3, 2014.

S. Meyer. A data structure for circuit net lists. In *Proceedings 25th Design Automation Conference*, pages 613–616. ACM/IEEE, June 1988.

S. Meyer. Adding methodological testing to naur's anti-formalism. In *IACAP 2013 Proceedings*. Internation Association for Computing and Philosophy, July 2013. URL `http://www.tdl.com/~smeyer/docs/naur-plus-method-300313-with-bohm.pdf`.

R. Morgan. *Building an Optimizing Compiler*. Digital Press, 1998.

P. Naur. *Computing as science*, pages Appendix 2, 208–217. naur.com Publishing, 2005. `www.Naur.com/Nauranat-ref.html` URL of Nov. 2015.

P. Naur. Computing versus human thinking. *Communication of the ACM*, 50(1):85–94, 2007.

P. Naur. *Conversations - Pluralism in Software Engineering*. Lonely Scholar Publishing, Belgium, 2011. Edgar Daylight ed.

S. Sutherland, S. Davidmann, and P. Flake. *System Verilog for Design*. Springer, 2006.

Tachyon-Design-Automation-Corp. Web page with *Signup and Download* link to click on to download cvc, October 2015. URL `http://www.tachyon-da.com`.

T. Thomas and P. Moorby. *The Verilog Hardware Description Language*. Springer, 2002.

US-Federal-Court. Pragmatic c software corp. v. antrim design sys., inc. 2003 u.s. dist. lexis 1824 (d. minn. 2003), 2003. URL `http://mn.findacase.com/research/wfrmDocViewer.aspx/xq/fac.20030128_0000185.DMN.htm/qx`. Legal precedent showing quality of CVC.

N. Wirth. The programming language pascal. *Acta Informatica*, 1:35–63, 1975.